arXiv:1401.6803v2 [cs.NI] 13 Oct 2014

On Efficiency and Validity of Previous Homeplug MAC Performance Analysis

C. Cano, D. Malone

Hamilton Institute
National University of Ireland, Maynooth
Co. Kildare, Ireland.

Abstract

The Medium Access Control protocol of Power Line Communication networks (defined in Homeplug and IEEE 1901 standards) has received relatively modest attention from the research community. As a consequence, there is only one analytic model that complies with the standardised MAC procedures and considers unsaturated conditions. We identify two important limitations of the existing analytic model: high computational expense and predicted results just prior to the predicted saturation point do not correspond to long-term network performance. In this work, we present a simplification of the previously defined analytic model of Homeplug MAC able to substantially reduce its complexity and demonstrate that the previous performance results just before predicted saturation correspond to a transitory phase. We determine that the causes of previous misprediction are common analytical assumptions and the potential occurrence of a transitory phase, that we show to be of extremely long duration under certain circumstances. We also provide techniques, both analytical and experimental, to correctly predict long-term behaviour and analyse the effect of specific Homeplug/IEEE 1901 features on the magnitude of misprediction errors.

Keywords: Power Line Communications, Homeplug, IEEE 1901, mean field analysis, decoupling approximation.

Email addresses: cristina.cano@nuim.ie (C. Cano), david.malone@nuim.ie (D. Malone)

1. Introduction

Data transmission using electrical wires, known as Power Line Communication (PLC), has the potential to become a useful complement and strong competitor to wireless networking solutions. The appeal of PLC relies in the high data rates it can deliver, its low deployment cost (as it is easy to retrofit) and allows communication through obstacles that commonly degrades wireless signals. Additionally, it also provides a low-cost alternative to complement existing technologies to reach ubiquitous coverage. For instance, as a backhaul for wireless sensor networks or small cells.

Research efforts have been mostly focused on the physical layer as the characteristics of PLC channels (including fading, impulsive noise and hidden/exposed terminal problems) impose several challenges on physical aspects [1]. However, they also have an impact on the Medium Access Control (MAC) protocol, which, in contrast, has not received much attention.

PLC standards (we focus on Homeplug [2] and IEEE 1901 [3]) define a MAC procedure similar to the Distributed Coordination Function (DCF) defined in the IEEE 802.11 standard for Wireless Local Area Networks (WLANs) [4]. PLC MAC protocols diverge from the vanilla DCF by adding a *deferral counter* that reduces the attempt rate when high contention is inferred on the channel (i.e., a certain number of packets are overheard). Despite being a substantial change over DCF, this feature has not been deeply evaluated. Attempts to compare DCF and PLC MAC protocols have been made in [5, 6, 7] and [8]. However, there is still much work to do in order to fully demonstrate the conditions under which the deferral counter improves the performance of the network.

Contributions to the analysis of PLC MAC protocols aim to extend the research on performance evaluation of PLC networks, crucial to advance understanding and optimisation. The only relatively complete analytic model of Homeplug MAC is presented in [9]. This analytic model has been widely used [10, 11, 12, 13, 14, 15], however we identify two significant issues. The first is the high computational complexity, making it unsuitable for online use or study of complex scenarios. Second, as we will demonstrate, the performance results obtained in the regime right before the predicted saturation point do not correspond to the long-term network behaviour. In detail, the main contributions of this work are the following:

- We propose a reformulation of the analytic model for the Homeplug MAC procedure presented in [9] that facilitates a simplified method

of solution. Our proposed analysis provides a 2 orders-of-magnitude improvement in runtime compared to the one presented in [9] while maintaining its accuracy. Specifically, we: *i)* apply a renewal reward approach [16, 17], *ii)* allow the most expensive operations to be precomputed and *iii)* provide and evaluate the accuracy of an optional exponential approximation to the probability to defer in a given backoff stage.

- We demonstrate that the results right before the predicted saturation point presented in [9] correspond to a transitory phase of the system instead to the long-term behaviour. We identify the two causes of misprediction in [9]: *i)* the decoupling approximation under infinite buffer size considered in the analytical model and *ii)* the presence, under certain circumstances, of an extremely long (of the order of magnitude of hours) transitory phase in experimental evaluation. This is the first work that shows the long duration of that transitory phase and, indeed, the causes of its occurrence suggest it may also be present in generic random access protocols. Having identified the causes of misprediction, we provide techniques to generate valid results from analytic models which use the decoupling approximation as well as from experimental studies based on simulations.

- We evaluate the impact of the deferral counter on the magnitude of misprediction errors by evaluating the different solutions obtained considering: *i)* the starting values of the deferral counter proposed by the standard, *ii)* no deferring (as done in IEEE 802.11 DCF) and *iii)* always deferring after overhearing following the proposal presented in [18].

The remainder of this article is organised as follows. In Section 2, we review the Homeplug MAC procedure. Next, in Section 3, related work on PLC MAC analytic models is discussed. Our simplified analytic model is described in Section 4. Then, in Section 5, we demonstrate that the previous performance results right before predicted saturation do not correspond to the long-term behaviour of the network. The performance evaluation, including the validation of the presented analytic model, an evaluation of the effect of the deferral counter on misprediction errors and a complexity comparison, are presented in Section 6. Finally, some final remarks are provided.

2. Homeplug MAC Background

The Homeplug MAC protocol is similar to the DCF defined in the IEEE 802.11 standard [4]. Each time a node has a new packet to transmit, the backoff stage ($i \in [1, m]$)[1] is initialised to 1 and a random backoff is selected among $[0, W_1]$. The backoff countdown is frozen when activity is detected on the channel and restarted when the medium becomes idle again. The packet is actually transmitted when the backoff countdown expires. If an acknowledgement is received, the packet is considered successfully transmitted. Otherwise, the node starts the retransmission procedure: the backoff stage changes to $i = \min(i + 1, m)$ and a new random backoff is selected among $[0, W_i]$, being W_i the contention window of stage i.

Additionally, in the Homeplug and IEEE 1901 MAC, a new counter, called the Deferral Counter (DC), is introduced. This counter is initialised at each backoff stage to M_i and decreased by one after overhearing a data packet or a collision. If a new packet or a collision are overheard and the value of the DC is equal to zero, the node acts as if a collision had happened: the backoff stage is increased if it has not yet reached its maximum value and a new backoff is selected among $[0, W_i]$. The goal of the DC is to avoid collisions when high contention is inferred by decreasing the aggressiveness of transmission attempts.

To provide channel access differentiation, four access categories (CA) are defined CA0–3. CA3 and CA2 share W_i and M_i values, as do CA1 and CA0 (see Table 1). Two slots (called PRS0 and PRS1) are allocated to allow nodes to announce the priority of their transmissions. The highest priority (CA3) is associated to both PRS0 and PRS1, the CA2 category is associated to PRS0 only, CA1 to PRS1 and the lowest access category (CA0) does not have any notification interval associated. Following this approach, stations know if there is a station with a frame that belongs to a higher category. In such a case, they postpone their transmissions until the high priority frames are released. However, in this work we consider the case where all stations contending for the channel belong to the same access category, as assumed in [9].

[1] Actually, ($i \in [0, m-1]$) but indexes have been relabelled here for clarity of illustration.

3. Related Work on Homeplug MAC Modeling

The analytic model in [9] takes into account all the features of the Homeplug MAC procedure except the channel access prioritisation. It models the node access procedure as a 3-dimensional Markov Chain in which the backoff and deferral counters, as well as the different backoff stages are considered. To solve for the stationary probabilities of the Markov Chain an iterative numerical method is required. Then, to compute the channel access delay, the authors follow a recursive approach, i.e., the delay at every backoff stage depends on the delay of previous stages. To obtain this performance metric, further iterations are needed. Moreover, to compute the channel access probabilities as well as the delay, many computationally expensive operations are required. The reason for this is the geometric nature of the deferral counter expiration and its dependence on the selected random backoff. The authors also extend the model to consider unsaturated conditions. However, the approach taken is based on further iterations in which the value of the transmission attempt rate is reweighted based on the offered load at every iteration until a reasonable approximation is obtained from the saturated model. As we will see in Section 6, the cost of solving this model makes it prohibitive for situations where results are required in real-time.

A similar approach to ours which also uses the renewal reward approach and that considers the backoff and deferral counters has been presented in [19]. However, only saturated conditions are considered. Recently, in [20], the same authors identify a different inaccuracy arising from the decoupling approximation under saturation conditions and a small number of contending stations, due to the high coupling between queues when the deferral counter is used. They also propose a model of the coupled system of queues based on the saturated assumption. Note that our work differs from [20] in the following ways: *i)* we evaluate the implications of the decoupling approximation on the buffers' occupancy, which can result in two different behaviours (see Section 5), and *ii)* we propose techniques to obtain the long-term performance of the network when using a model based on the decoupling approximation under unsaturated conditions. We note that directly extending the analysis of the coupled queues to nonsaturated conditions appears intractable.

Two other analytic models of Homeplug exist in the literature ([21] and [18]). However, they model mechanisms that differ from the original standardised Homeplug procedure. In [21], the value of the contention window is fixed for all backoff stages, while in [18] the backoff stage is incremented

Parameter	All CAs	Parameter	CA3/2	CA1/0
M_1	0	W_1	7	7
M_2	1	W_2	15	15
M_3	3	W_3	15	31
M_4	15	W_4	31	63

Table 1: Parameters of the different access categories in Homeplug

every time a new packet is overheard (thus having $M_i = 0,\ \forall i$). These new approaches obviously lead to a simplified analysis, as one of the dimensions of the Markov Chain is removed. However, the flexibility the standard provides by allowing to tune both the W_i and M_i is not captured by these frameworks.

Consequently, the analytic model in [9] has been widely used both to derive performance metrics and as a basis for extension [10, 11, 12, 13, 14, 15] as it strictly follows the procedure defined in the standard and it considers saturated as well as unsaturated conditions. This widespread use motivates our improvements in terms of *i)* reducing complexity while maintaining the accuracy and retaining all features defined in the standard [2] and *ii)* understanding the transitory nature of the results predicted before saturation in [9].

4. Simplified Homeplug Analytical Model

In this section, we present the simplified analytic model of the Homeplug MAC protocol. To this end we take a renewal reward approach [16, 17] motivated by the fact that the attempt rate of a given node can be viewed as a regenerative process. This approach allows us to compute metrics of interest without the need to solve all state probabilities of the Markov Chain embedded in the analysis. Similar approaches have already been proposed to simplify the derivation of analytic models of the IEEE 802.11 for WLANs, see for instance [22] and [23].

We also apply the decoupling approximation to model the conditional (given that a packet is transmitted) collision probability independently of the backoff stage, as also done in [9] and in a large number of IEEE 802.11 analytic models (see [24] for a complete survey on this topic). Although this assumption may not be valid [25], it allows us to simplify the analysis while still accurately predicting the metrics of interest in a broad range of circumstances. The decoupling approximation is also used to model the

buffer occupancy probability after a transmission as being independent of the backoff stage at which the packet was transmitted. This assumption is implicitly made in [9]. The accuracy of this assumption will be discussed in detail in Section 5 as it is contributing factor to misprediction of results right before saturation in [9].

The rest of assumptions and considerations taken into account are: *i)* an infinite, or large enough to be considered infinite, buffer size and retry limit, *ii)* exponentially distributed interarrival of packets, *iii)* ideal channel conditions (i.e., packets are always received correctly in case of no collision), *iv)* contention among a single access category and *v)* all nodes are in mutual coverage range, that is, all nodes can overhear each other's transmissions. These are the same assumptions considered in [9] and in a number of IEEE 802.11 analytical frameworks. On one hand they allow us to compare the accuracy of our analysis and the one in [9] under the same conditions. On the other hand, they are useful to assess and understand the performance of the MAC procedure independently of other factors that may have an influence in the results in practice.

4.1. Renewal Reward Approach

Assuming an infinite buffer size, the mean queue occupancy (ρ) of a node is derived by considering the time needed to release a packet from the queue (X), called service time or MAC access delay, and the packet arrival rate from the network layer (λ) as $\rho = \min(\lambda X, 1)$. While λ depends on the application, the service time is computed as the sum of the following three components: *i)* the total average backoff duration until the successful frame transmission, *ii)* the total time on average spent in transmitting packets that result in a collision and *iii)* the time spent successfully transmitting the packet:

$$X = E[w]\alpha + (n_{\mathrm{t}} - 1)T_{\mathrm{c}} + T_{\mathrm{s}}, \tag{1}$$

where $E[w]$ denotes the average number of slots in backoff, α is the average slot duration and n_{t} is the average number of attempts to successfully transmit a packet. The duration of a successful transmission (T_{s}) and a collision (T_{c}) are computed as shown in Eq. 2. These durations account for the transmission notification intervals (PRS0 and PRS1), the time to transmit the actual frame (T_{fra}) and the acknowledgement (T_{res}), as well as the

contention and response interframe spaces (CIFS and RIFS)[2], see [2].

$$T_{\mathrm{s}} = T_{\mathrm{c}} = \mathrm{PRS0} + \mathrm{PRS1} + T_{\mathrm{fra}} + \mathrm{RIFS} + T_{\mathrm{res}} + \mathrm{CIFS} \tag{2}$$

Under the decoupling assumption with an infinite number of retries, the average number of attempts to transmit a frame (n_{t}) is computed as shown in Eq. 3.

$$n_{\mathrm{t}} = \frac{1}{1-p}, \tag{3}$$

where the conditional collision probability (p) is obtained as the complementary of having at least one of the other $n-1$ nodes transmitting a frame in the same slot (Eq. 4), with τ denoting the attempt rate of a node.

$$p = 1 - (1-\tau)^{n-1} \tag{4}$$

We view the attempt rate as a regenerative process, where the renewal events are when the MAC begins processing a new frame. Thus, we apply the renewal reward theorem (Eq. 5).

$$\tau = \frac{n_{\mathrm{t}}}{E[w] + n_{\mathrm{t}} + I} \tag{5}$$

The term I in Eq. 5 accounts for the number of slots in idle state (when there is no packet waiting in the queue for transmission) and is computed as the probability of having an empty queue over the probability of a packet arrival in a slot. Considering an M/M/1 queue, we then compute I as in Eq. 6.

$$I = \max\left(\frac{1-\rho}{1-e^{-\lambda\alpha}}, 0\right) \tag{6}$$

The average slot duration while the node is in backoff is derived depending on the type of slot that is overheard (Eq. 7). A slot can be empty if no other node transmits (that occurs with p_{e} probability) and, in such a case,

[2]In contrast to [9], we have considered $T_{\mathrm{s}} = T_{\mathrm{c}}$ to closely model the standard procedure. Observe that, in case of collision or frame errors, either an NACK is received or nodes wait an EIFS to provide protection from collisions for ongoing transmissions [2]. Thus, the duration of a collision can be better approximated to that of a successful transmission.

$$E[w_i] = \frac{1}{W_i+1}\left[\frac{M_i^2+M_i}{2} + \sum_{k=1}^{W_i-M_i}\left\{\sum_{l=0}^{k-1}\binom{M_i+l}{l}(1-p_\mathrm{b})^l p_\mathrm{b}^{M_i} p_\mathrm{b}(l+M_i+1) + \sum_{l=0}^{M_i}\binom{M_i+k}{k+l}(1-p_\mathrm{b})^{k+l} p_\mathrm{b}^{M_i-l}(k+M_i)\right\}\right] \tag{9}$$

its duration is σ. Otherwise, it can be occupied due to a successful transmission (that happens with probability p_s) or a collision (that occurs with p_c probability), with durations T_s and T_c, respectively.

$$\alpha = p_\mathrm{s}T_\mathrm{s} + p_\mathrm{c}T_\mathrm{c} + p_\mathrm{e}\sigma \tag{7}$$

Probabilities p_s, p_e and p_c are obtained as follows:

$$\begin{aligned} p_\mathrm{s} &= (n-1)\tau(1-\tau)^{n-2}, \\ p_\mathrm{e} &= (1-\tau)^{n-1}, \\ p_\mathrm{c} &= 1 - p_\mathrm{s} - p_\mathrm{e}. \end{aligned} \tag{8}$$

To compute X, it remains to obtain the value of the average number of slots spent in backoff until a packet is successfully transmitted ($E[w]$). With this aim, we follow the same approach as in [9] but we compute the average number of slots at every backoff stage. The average number of slots waiting in backoff at backoff stage i (with $1 \leq i \leq m$) is shown in Eq. 9. This expression is obtained from [9] but it has been rearranged to account for the number of backoff slots instead of the total delay to transmit a packet. The reader is referred to [9] for details.

In Eq. 9, p_b refers to the probability of overhearing another transmission while the node is in backoff:

$$p_\mathrm{b} = 1 - (1-\tau)^{n-1}. \tag{10}$$

A node moves to the next backoff stage (if it has not reached its maximum value) whenever it fails to transmit a packet. We define a failure here as the case in which either: *i)* the backoff expires leading to the transmission of

the packet and there is a collision or *ii)* the deferral counter expires so that the node defers its transmission. The probabilities of each of these situations depend on the backoff stage as they are a function of the value of W_i and M_i. Thus, we compute the probability of failure at backoff stage i as shown in Eq. 11.

$$p_{\mathrm{f}}^{(i)} = p \cdot p_{\mathrm{bo}}^{(i)} + p_{\mathrm{defer}}^{(i)}, \tag{11}$$

where $p_{\mathrm{bo}}^{(i)}$ is the probability that the backoff expires in backoff stage i. On the other side, $p_{\mathrm{defer}}^{(i)}$ denotes the probability that the deferral counter reaches zero and a new packet is overheard in the given backoff stage. This probability is also obtained from [9] and it is calculated as:

$$p_{\mathrm{defer}}^{(i)} = \frac{1}{W_i + 1} \sum_{k=1}^{W_i - M_i} \sum_{l=0}^{k-1} \binom{M_i + l}{l} (1 - p_{\mathrm{b}})^l p_{\mathrm{b}}^{M_i} p_{\mathrm{b}}. \tag{12}$$

The probability that the backoff expires is just the complementary of the probability to defer (Eq. 13).

$$p_{\mathrm{bo}}^{(i)} = 1 - p_{\mathrm{defer}}^{(i)} \tag{13}$$

Thus, the total average number of slots in backoff is obtained by summing the number of slots waiting at each stage weighted by the probability of moving to that backoff stage (see Eq. 14). Note that the probability to move to a given backoff stage is the probability to face a failure in the previous ones. An extra term is considered for the last backoff stage as the node returns to it after every failure. Assuming a packet is retransmitted until it is successfully transmitted the last term follows.

$$\begin{aligned} E[w] = E[w_1] + \sum_{i=2}^{m-1} E[w_i] \prod_{j=1}^{i-1} p_{\mathrm{f}}^{(j)} \\ + E[w_m] \prod_{j=1}^{m-1} p_{\mathrm{f}}^{(j)} \frac{1}{1 - p_{\mathrm{f}}^{(m)}} \end{aligned} \tag{14}$$

Finally, having the payload length (L), we obtain the throughput as:

$$S = \rho \frac{L}{X} \tag{15}$$

The previous expressions can be solved by using an iterative numerical method. If compared to the previous analytic model in [9], the renewal reward approach allow us to reduce the complexity by removing two iteration loops: the one involving the calculation of the state probabilities of the Markov Chain as well as the one regarding the total average MAC access delay computation.

4.2. Precomputation of $p_{\text{defer}}^{(i)}$ and $E[w_i]$

The most computationally expensive operations of our simplified analytic model are the computation of $p_{\text{defer}}^{(i)}$ and $E[w_i]$ (Eq. 9 and 12, respectively). However, both are a function of just three variables: $p_{\text{defer}}^{(i)} = f(W_i, M_i, p_{\text{b}})$ and $E[w_i] = g(W_i, M_i, p_{\text{b}})$. Since W_i and M_i are fixed parameters and $p_{\text{b}} \in [0, 1]$, we can easily precompute the values of these two metrics for each relevant value of W_i, M_i and p_{b}.

The error of the approximated results can be reduced by applying a linear interpolation between the two closest precomputed values of p_{b}. Results presented in this work have been obtained using the precomputed values as they have been found indistinguishable (using a step size equal to 10^{-4} for p_{b}) from the ones obtained using Eq. 9 and 12.

4.3. Exponential Approximation to $p_{\text{defer}}^{(i)}$

We consider a further optional step in reducing the complexity of the analytic model. We may calculate the probability to defer a packet in a given backoff stage by approximating both the backoff countdown and the deferral counter with exponential distributions with means chosen to match the actual distributions. We approximate the backoff counter and the deferral counter exponentially distributed with rate $\beta_i = 2/W_i$ and $\gamma_i = p_{\text{b}}/(M_i + 1)$, respectively. Then, we can easily derive the probability to defer as the probability that the deferral counter expires first, that is:

$$p_{\text{defer}}^{(i)} = \frac{\gamma_i}{\beta_i + \gamma_i}. \tag{16}$$

Although the otional approximation considerably reduces the complexity of the calculation of $p_{\text{defer}}^{(i)}$ in Eq. 12, it only provides a rough estimate as shown in Fig. 1(a) (CA3/2) and 1(b) (CA1/0) for an arbitrary value of p_{b}. The disagreement is more notable for high values of M_i. The reason for this result is the inability of the approximation to model the cases in which the random backoff selected at stage i is smaller than M_i. In such a case,

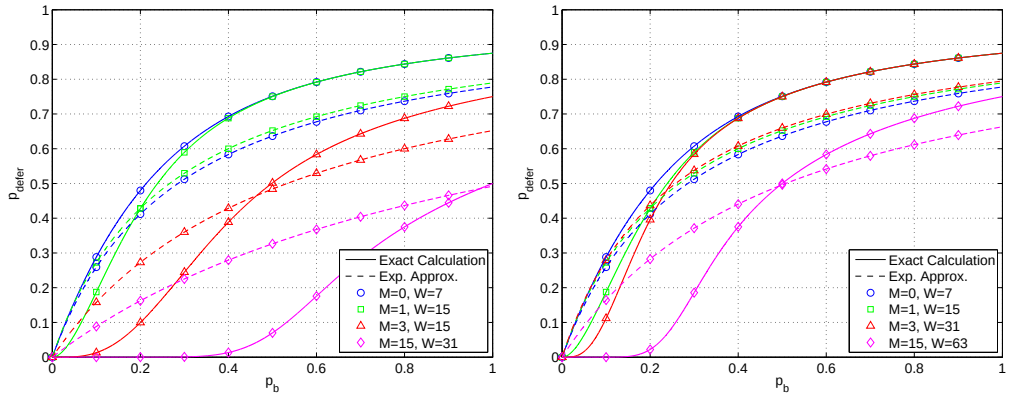

(a) CA3/2 (b) CA1/0

Figure 1: Comparison of $p_{\text{defer}}^{(i)}$ computed following the exact calculation and the exponential approximation.

regardless of the value of p_{b}, the backoff counter will always expire first, yet the exponential approximation always gives some probability that either expire first. However, note that the estimation is accurate for certain values of p_{b}, M_i and W_i. As will be shown in Section 6, the approximation does provide a good estimate of throughput and delay for certain ranges of p and thus, could be safely used in the cases identified.

5. On the Performance Results Right Before Saturation

Before the performance evaluation, in this section, we demonstrate that the results presented in [9] just prior to predicted saturation correspond to a transitory phase of the network. To do so we refer to previous literature regarding the decoupling approximation and to show that the coupled system of queues under infinite buffer size is unstable in this regime. Thus, with an infinite queue size, higher throughput than the one found in saturation, as predicted in [9], cannot be maintained in the long term. Furthermore, we support this demonstration by experimental evaluation in which we show, under certain circumstances, the extremely long time during which the network remains in the transitory period before moving to the long-term behaviour.

5.1. The Decoupling Approximation

The decoupling approximation allows to make the analysis tractable by assuming that each queue can be modelled independently of the dynamics

of the rest of queues in the network. However, one has to be careful when interpreting the results obtained. As already observed in [24], unsaturated analytic models that decouple the queue dynamics can provide two different solutions in certain regimes. In particular, when the packet arrival rate is slightly higher than the maximum load that the system could serve in saturation. These analytic models do not consider the number of instantaneous contending stations, i.e., the number of queues that have at least a packet buffered at the same time. Neglecting this fact makes it impossible to identify the actual regime at which the queues operate in this region. This is caused by the possibility of facing two extreme cases: the queues being mostly empty or saturated conditions. On one hand, if a reduced number of nodes are contending for the channel, the channel access delay is small, thus, making unlikely that a high number of packets accumulate in the buffer for transmission. This implies higher throughput than found in saturation, as the conditional collision probability is reduced. On the other hand, if a higher number of nodes are contending, the channel access delay increases and so does the probability of having more packets accumulating for transmission. The effect of the latter case, under infinite buffer size, is the saturation of the buffer. Analytical models that do not consider the number of instantaneous contending stations are not always able to differentiate among the two different solutions and thus, iterative algorithms may converge to either, depending on the starting parameters.

5.1.1. Which solution corresponds to the performance of the network in the long term?

As already conjectured in [24], when considering infinite buffer size, the lowest throughput solution is the one that corresponds to the long-term behaviour of the network, while the other solution is just transitory. This has been demonstrated in [26], where a coupled system of parallel queues with infinite buffer size and state-dependent service rates is analysed. The system of queues is found to be unstable (the stability limit is surpassed) when the packet arrival rate (λ) is higher than the service rate that can be achieved when all nodes are simultaneously contending for the channel. This service rate corresponds to the one found in saturation. Therefore, when the packet arrival rate is above the service rate in saturation, all queues will eventually have an increasing number of packets. Thus, the long-term performance of the network corresponds to that in saturation (all nodes will have at least a packet buffered for transmission). The results presented in [9] show

higher throughput than the one found in saturation when the stability limit is slightly surpassed. However, given that the queues are unstable in that regime, the long-term throughput should correspond to the one in saturation instead.

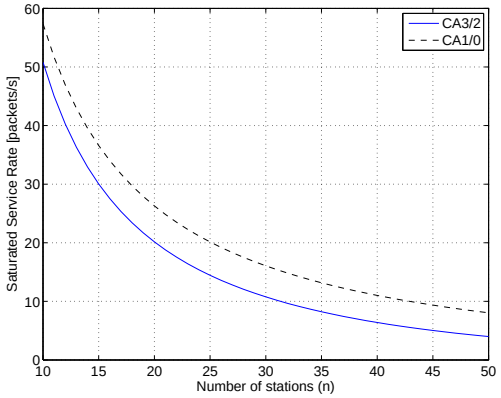

Figure 2: Service rate in saturated conditions.

In order to identify the stability limit regime, Fig. 2 shows, for different number of nodes contending for the channel and access categories, the service rate obtained in saturation conditions. The parameters used are shown in Table 2 and correspond to those in [9], which are defined in Homeplug 1.0 [2]. Note that the service rate in saturated conditions is just $\mu_{\text{sat}} = 1/X_{\text{sat}}$, where X_{sat} denotes here the service time considering that all nodes always have a packet ready for transmission. From [26] (for a homogeneous system), we identify the stability condition of the system of coupled queues as $\lambda < \mu_{\text{sat}}$. As will be shown in Section 6, when the stability condition of the network is not satisfied, there is a potential for obtaining two solutions from the analytic model. The stability limit is surpassed if λ is higher than the service rate depicted. It is important to emphasise that the potential for obtaining two solutions decreases when $\lambda >> \mu_{\text{sat}}$ as the probability of having a large number of the nodes simultaneously contending for the channel increases.

When using iterative numerical solvers, setting certain starting parameters can provide a particular solution among the two possible ones. Setting the metrics assuming a highly-loaded network starts the numerics closer to the saturated solution. On the contrary, assuming lightly-loaded conditions, starts the iterative loop closer to the transitory phase. When there is only one solution, the solver converges to it independently of starting under any

of these initial conditions. We will use this technique in Section 6, to obtain both solutions from the analysis. This technique, although useful to distinguish the two solutions, is not able to differentiate them when both are very close. However, we know that after surpassing the stability limit (as described above) the throughput that should be obtained is the one corresponding to saturation.

5.2. The Long Transitory Phase in Experimental Evaluation

When $\lambda >> \mu_{\text{sat}}$, the system rapidly moves to the long-term behaviour. However, performing an experimental evaluation right after the stability limit can provide wrong results since, as we have found in this work, the length of the transitory phase can be extremely long. We have found that if the experiments are started with the queues empty, it can take a long time to reach the long-term performance since the system has to reach a point at which a large number of nodes are simultaneously contending for the channel and start to have their queues filled with an increasing number of packets. The possibility of a long transitory phase for random MAC protocols in this regime was postulated in [27], however no experimental findings or proof was provided. Here we take the step of showing experimentally that for Homeplug MAC this is exactly the case.

For that purpose we have run long simulations with load just above the stability limit with queues long enough to be considered infinite (1000 packets). From the simulations we have measured the throughput and the queue size (maximum, average and minimum) of all the nodes in the network in 1s-long time intervals. Simulation results are obtained using a custom simulator based on the SENSE framework [28]. We have considered the case in which 50 nodes with packets belonging to CA3/2 access categories are accessing the channel, as larger numbers of nodes help us highlight the discrepancy between the previously presented results (see [9] Fig. 6 and 7) and the saturated solution. We begin the simulations with the queues empty.

Observe in Fig. 3(a,c) that, for different values of the packet arrival rate, there is a sharp change from the transitory behaviour to the long-term solution. In Fig. 3(b,d), we zoom in on the queue behaviour around this sharp change. We see that the queues start to fill with packets until the maximum size is reached. The conjecture that the transitory phase is of long duration is also supported, as we observe a duration of the order of hours in the runs shown.

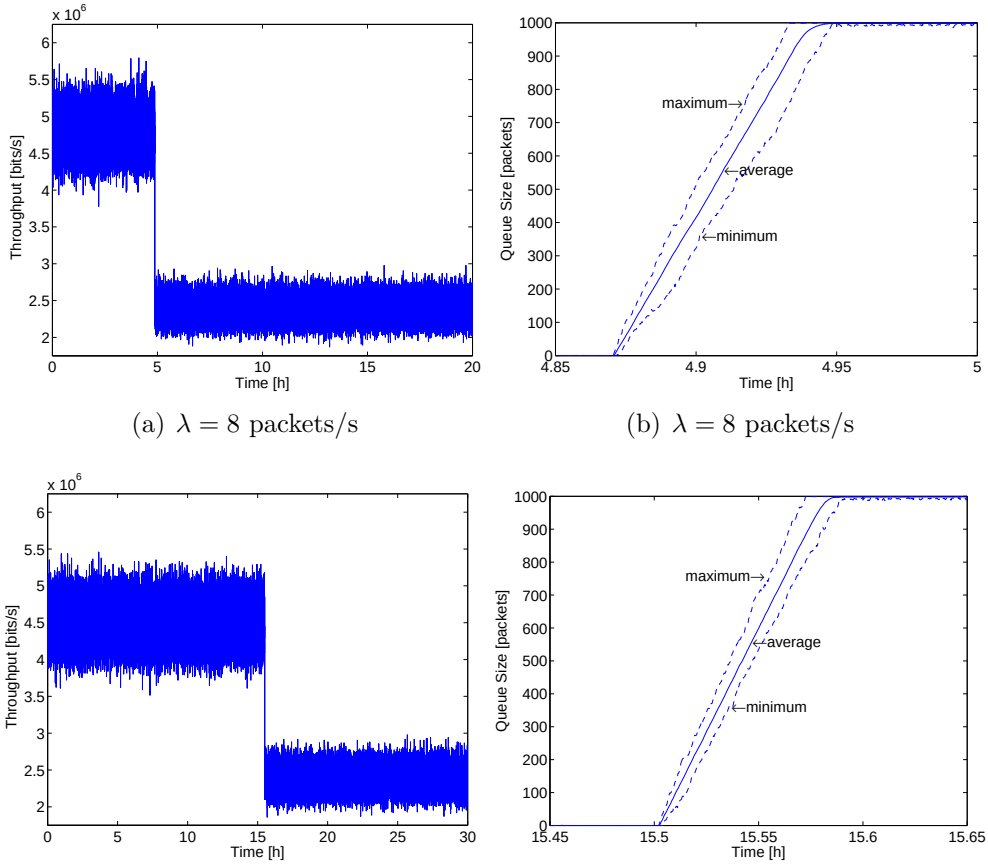

(a) $\lambda = 8$ packets/s (b) $\lambda = 8$ packets/s

(c) $\lambda = 7.5$ packets/s (d) $\lambda = 7.5$ packets/s

Figure 3: Evolution of throughput and queue size (CA3/2, $n = 50$ nodes).

It is important to emphasise that the results presented in Fig. 3 are each from a single simulation run. The time at which the change in behaviour occurs varies in different runs due to its stochastic nature. To characterise the length of the transitory period one must perform analysis considering the coupled dynamics of the system of queues. Due to its difficulty, especially considering the complexity of the Homeplug MAC access procedure, we consider this analysis out of the scope of this work.

Because of this long transitory period, care is required in designing the simulations. One way to obtain the long-term performance is to start with the queues empty, run the simulation for a long time until the queues are saturated and then start taking the statistics of the performance metrics of

interest. However, we suggest starting the simulations with a number of packets preloaded in the queues as a more practical way to force the system to enter into the long-term operating state. If the queues are stable, nodes will be able to release these preloaded packets in a reasonable amount of time. If the queues are unstable, we have started the simulation closer to their stationary regime, and will see the long-term throughput more quickly. We use this technique to avoid the transitory period in the next section in order to suppress the effect of the transitory on performance results.

6. Performance Evaluation

In this section, we present the validation of our analytic model. The validation is divided according to saturated and unsaturated conditions. The later also show the two different solutions the analytic model may provide as well as the long-term performance to which the system converges to, obtained from experimental evaluation. We then evaluate the effect of the deferral counter on misprediction errors and, in the last subsection, we show the complexity reduction of our simplified analysis.

Simulations are performed, as in the last section, using the SENSE framework and parameters shown in Table 2. Results show averages from simulation runs of 10000 s.

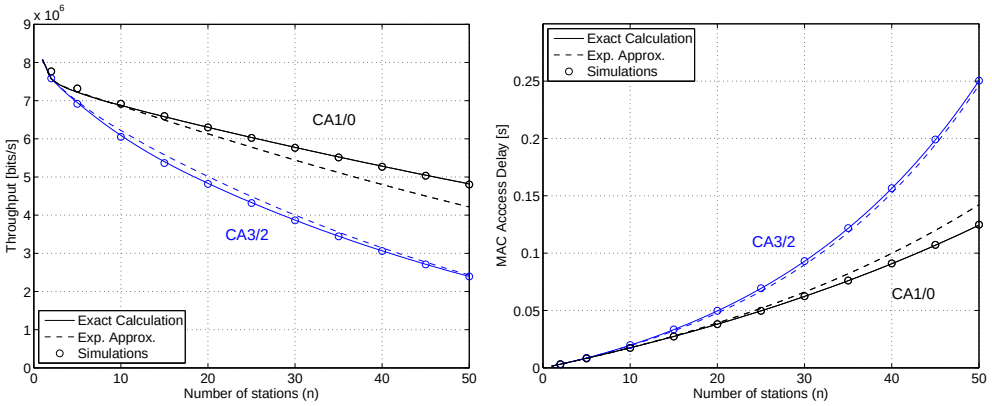

(a) Throughput (b) MAC access delay

Figure 4: Performance results in saturation conditions. Comparison among the exact calculation, the exponential approximation of $p_{\text{defer}}^{(i)}$ and simulations.

Parameter	Value in Homeplug 1.0
Data rate (R)	14 Mbps
Frame transmission time (T_{fra})/L	1153.5 μs/1500 bytes
ACK transmission time (T_{res})	72 μs
Slot time (σ)	35.84 μs
Data-ACK interframe space (RIFS)	26 μs
Contention interframe space (CIFS)	35.84 μs
Tx. indication slots (PRS0 = PRS1)	35.84 μs

Table 2: Parameters Homeplug 1.0

6.1. Saturated Conditions

Throughput and channel access delay metrics in saturation conditions for different number of nodes are shown in Fig. 4(a) and 4(b), respectively. Results from *i)* the exact computation (obtained from precomputing the values of $p_{\text{defer}}^{(i)}$ and $E[w_i]$), *ii)* the exponential approximation to $p_{\text{defer}}^{(i)}$ and *iii)* simulations are depicted. Observe that the results obtained from the exact calculation are in agreement with the performance predictions and simulations presented in [9] in Figs. 4 and 5 (slight differences appear due to our $T_{\text{s}} = T_{\text{c}}$ consideration). Moreover, we also see how the optional approximation is able to accurately predict both throughput and channel access delay in all the cases considered for access categories CA3/2. For access categories CA1/0, the exponential approximation is accurate for a small number of contending nodes (less than 15). For a large number of nodes, the approximation provides a rough estimate.

To study when the exponential approximation is useful, Fig. 5 shows the channel access delay vs. the conditional collision probability, obtained using the exact calculation and the exponential approximation. This plot allows us to identify when the exponential approximation gives a good estimate, independent of the exact configuration, since the channel access delay as a function of the conditional collision probability is independent of the number of nodes contending [29]. Thus, we observe that the exponential approximation is accurate for access categories CA3/2 for all configurations, while for access categories CA1/0 it provides a good estimate when $p < 0.4$.

6.2. Unsaturated Conditions

Results in unsaturated conditions for $n = 10, 30$ and 50 and varying the packet arrival rate λ are shown in Fig. 6. Values of throughput and

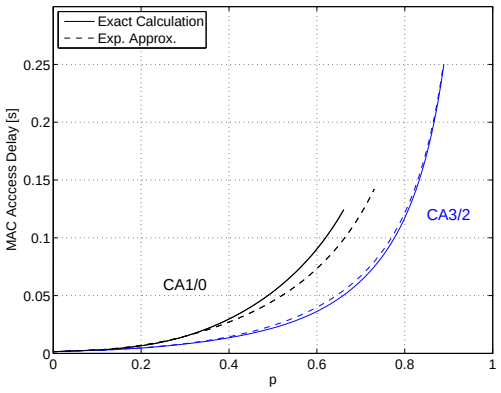

Figure 5: Channel access delay vs. p.

channel access delay for the different access categories obtained from the exact computation and simulations are depicted.

For the analytic results, we show two solutions, corresponding to beginning the numerics for solving the model with $I = 0$ slots (no idle periods between a departure and the next packet arrival, labelled *Analysis 1*) or $I = 1000$ slots (long idle periods between departures and next packet arrivals, labelled *Analysis 2*). Observe in Fig. 6 that the model may converge to different solutions. As expected, just after the stability limit, the analytic model provides two different solutions depending on the starting values we provide to the iterative loop. We see that *Analysis 1* provides the lowest throughput, while *Analysis 2* shows higher throughput.

We use here the technique of preloading the queues as already discussed in the last section. Thus, we have averaged the values of 10000 s-long simulations with the queues preloaded with 50 packets. We have compared results from *Analysis 1* and find agreement with the long-term performance obtained after the transitory phase.

Note how, excluding the results right after saturation (as we already discussed), the results obtained are also in agreement with the outcomes obtained in [9] in Figs. 6-9, validating, as in the saturated case, the accuracy of our proposed analysis.

Finally, we compare the optional exponential approximation and the exact calculation (both starting with $I = 0$ slots) in unsaturated conditions in Fig. 7. When the system is saturated, we obtain the accuracy presented in the last subsection. Here it can be observed how the exponential approxi-

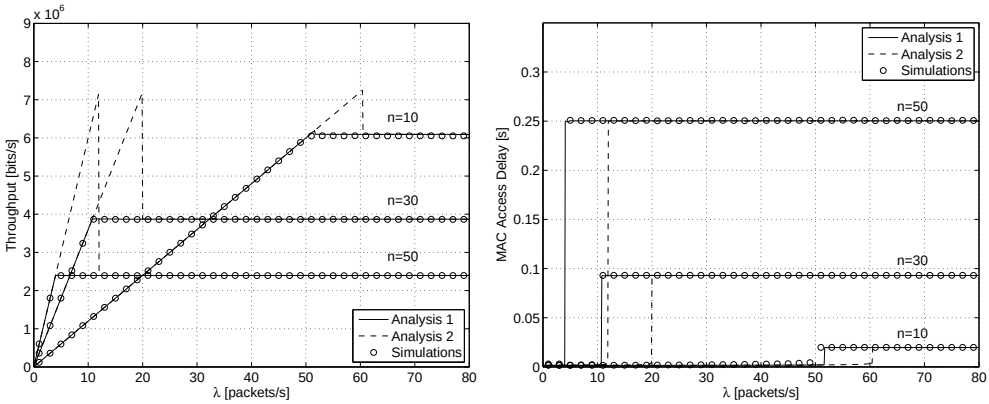

(a) Throughput CA3/2 (b) MAC access delay CA3/2

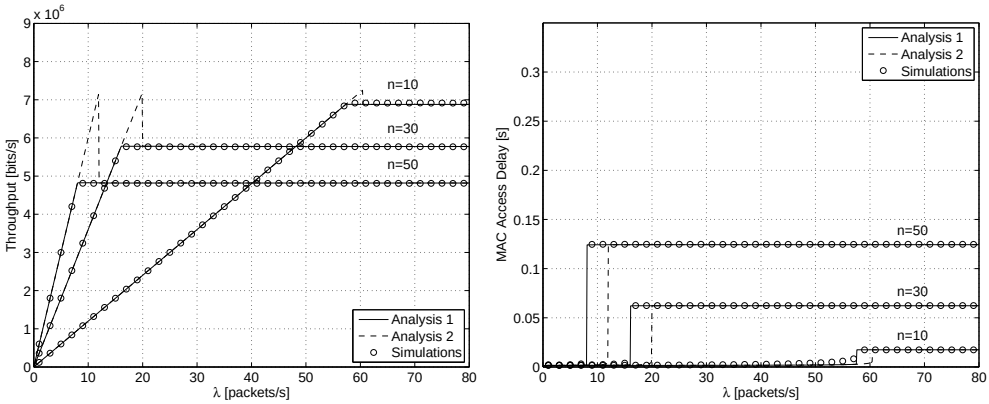

(c) Throughput CA1/0 (d) MAC access delay CA1/0

Figure 6: Performance results in unsaturated conditions. Comparison among the two solutions derived from the exact analysis and simulations.

mation is able to predict the saturation point with a difference, for all cases evaluated, smaller than $\Delta\lambda = 2$ packets/s.

6.3. Effect of the Deferral Counter

We now evaluate the effect of the the deferral counter on the difference between the two different solutions obtained from the analytic model right after the stability limit. These results allow us to quantify the potential misprediction errors. We consider *i)* the results with no deferral (i.e., corresponding to a DCF-like access procedure [4]) and *ii)* increasing the backoff

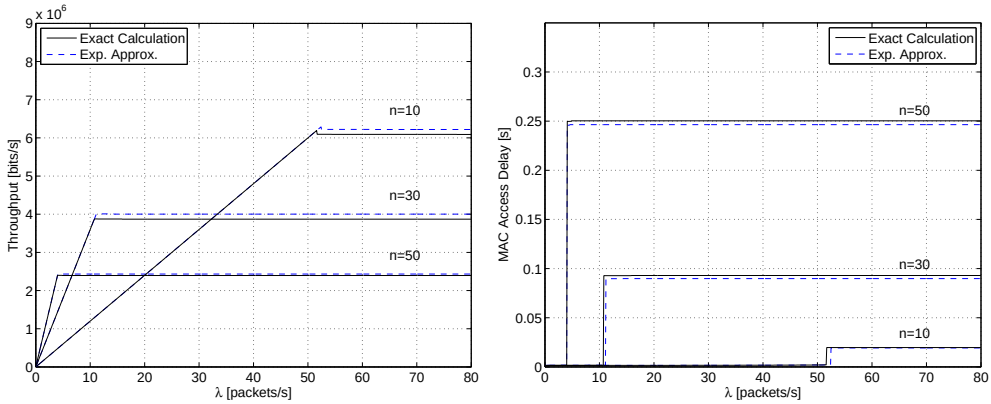

(a) Throughput CA3/2

(b) MAC access delay CA3/2

(c) Throughput CA1/0

(d) MAC access delay CA1/0

Figure 7: Performance results in unsaturated conditions. Comparison among the exact analysis and the exponential approximation to $p_{\text{defer}}^{(i)}$.

stage every time a packet from a neighbouring node is overheard (i.e., the proposal presented in [18] for improving the performance of Homeplug MAC).

6.3.1. No deferral

To obtain the results without the deferral counter (i.e., $M_i = \infty,\ \forall i$), we run the analytic model described in Section 4 with $p_{\text{defer}}^{(i)} = 0, \forall i$ and $E[w_i] = (W_i + 1)/2$. Note that the performance obtained will match that of DCF with the same parameters. However, for comparison purposes, we use the MAC parameters of Homeplug MAC shown in Table 2. We start

the numerics for solving the model as in the last subsection: *Analysis 1* corresponds to an initial value of $I = 0$ slots and *Analysis 2* corresponds to the results starting with $I = 1000$ slots. The two different solutions for throughput obtained from the analytic model are shown in Fig. 8. Note that compared to the results depicted in Fig. 6, the difference between the two solutions provided by the analytical model is higher in Fig. 8. The saturation throughput obtained with no deferral is smaller compared to the one obtained using the deferral counter. Thus, with no deferral (Fig. 8), the saturation throughput (*Analysis 1*) and the throughput obtained under a lower contention assumption (*Analysis 2*) show a bigger difference than the one found with deferral (Fig. 6). Thus, under the same configuration, protocols such as the DCF have the potential to be affected by higher misprediction errors than Homeplug MAC, especially when small values of W are used as in the case of CA3/2 categories.

6.3.2. Deferring always after overhearing

In accordance with the conclusions drawn for no deferral, we expect that deferring always after overhearing will show a smaller difference among the two solutions as this setting increases the saturation throughput with higher contending nodes, especially when considering an increased number of backoff stages and larger contention windows. The improvement in the saturation throughput by setting $M_i = 0, \ \forall i$ has already been studied for different configurations of i, W and payload lengths (see Fig. 3 in [18]). However, in order to perform a direct comparison with the outcomes presented in previous sections, we show in Fig. 9 the results using the same parameters considered throughout this work (i.e., those shown in Table 1 and 2) instead of referring to results in [18]. Moreover, this configuration allows us to evaluate the benefit of increasing the backoff stage after every packet overheard when the number of backoff stages are set to those recommended by the standard.

Observe that, under these particular settings, the saturation throughput compared to Fig. 6 is substantially improved but the potential misprediction errors are still considerable.

6.4. Reduction of Complexity

Here, we evaluate the reduction in complexity obtained using our reformulated analysis (with precomputed $p_{\text{defer}}^{(i)}$ and $E[w_i]$) and also our optional exponential approximation in comparison to the analysis presented in [9]. Moreover, we also compare the time needed to run the analytic models and

n	Analysis in [9]	Exact Analysis	Exp. Approx.	Simulations
10	584.5 s	3.7 s	1.7 s	165.5 s
50	420.0 s	4.2 s	3.5 s	866.2 s

Table 3: Elapsed Runtime Comparison for 10/50 nodes

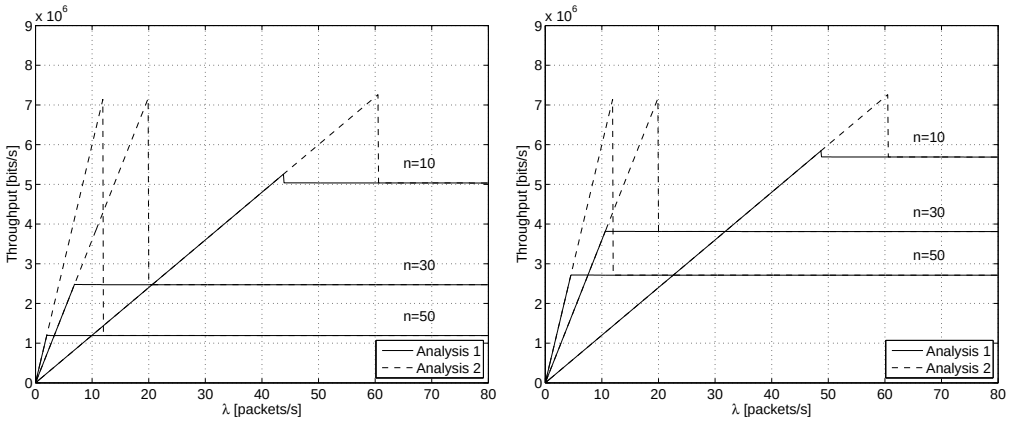

(a) Throughput CA3/2

(b) Throughput CA1/0

Figure 8: Throughput in unsaturated conditions with no deferral ($M_i = \infty$, $\forall i$). Comparison among the two solutions derived from the exact analysis.

10000 s-long simulations, as a crucial value of analysis is as a fast way to predict performance metrics. Since the analytic models are solved by means of iterative numerical methods, a regular complexity analysis based on the number of operations is not useful as the number of iterations needed may vary and cannot be predicted in advance. For this reason, we have obtained the elapsed times to run the analytic models (using *tic* and *toc* commands in Matlab R2011b) and simulations (using the *time* command in Ubuntu). Tests have been performed in an Intel Core i5 at 3 GHz with 6 GB of RAM running a 12.04 32-bit Ubuntu operating system. Results of a single run considering 10 and 50 nodes using access categories CA1/0 and in saturated conditions are depicted in Table 3. Obviously, these results may be affected by implementation issues. However, as can be observed, the improvement of our simplified analysis (both the exact analysis and the exponential approximation) is 2 orders of magnitude. Furthermore, the analysis in [9] has been shown to be of a similar computational expense to simulations, placing reasonable doubts on its efficiency.

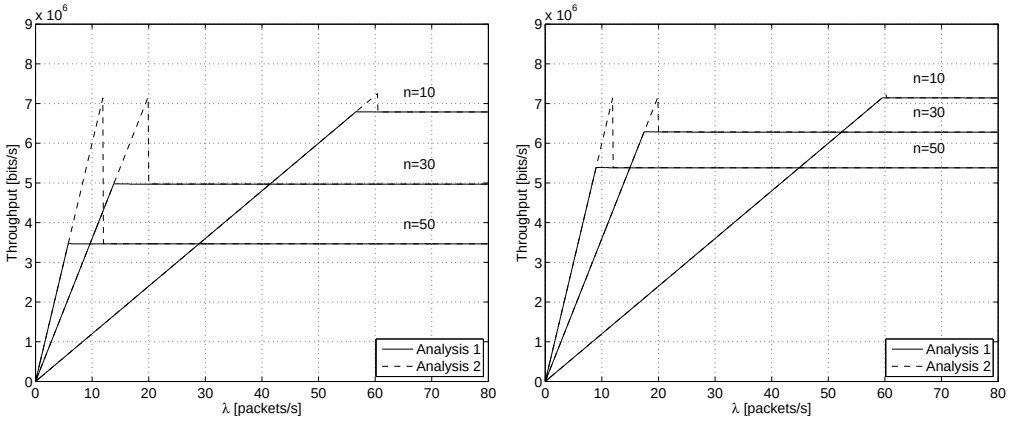

(a) Throughput CA3/2 (b) Throughput CA1/0

Figure 9: Throughput in unsaturated conditions deferring always after overhearing ($M_i = 0$, $\forall i$). Comparison among the two solutions derived from the exact analysis.

7. Final Remarks

In this work we have reformulated the analytic model of the Homeplug MAC procedure presented in [9] while maintaining its accuracy. To that end, we have applied a renewal reward approach to model the channel attempt rate instead of solving for the state probabilities associated with the Markov chain. Furthermore, following our approach, the most computationally expensive operations can be precomputed. We have considerably reduced the complexity of the analysis, resulting in a two orders-of-magnitude improvement in runtime. An optional exponential approximation has also been proposed and shown to be generally accurate for access categories CA3/2 and accurate under a small-conditional-collision-probability condition for access categories CA1/0.

Building upon previous literature on mean field analysis of random access MAC protocols and on stability analysis of systems of parallel queues with state-dependent service rates, we have demonstrated that the results right before the predicted saturation point in [9] correspond to a higher-throughput transitory phase of the network. We have also experimentally supported this finding and shown the extremely long duration of the transitory period.

The reformulated analytic model has been validated in saturated and unsaturated conditions using simulations. Furthermore, we have found agreement (except right after the stability limit, as previously discussed) with the

results presented in [9], establishing the good accuracy of our proposal. In unsaturated conditions, we have shown the two solutions that the analytic model provides right after the stability limit and the long-term behaviour the network converges to by means of experimental evaluation. As expected, the long-term performance is in agreement with one of the solutions provided by the analytic model.

In this article we have presented the first long-term performance evaluation of Homeplug MAC under infinite (or large enough to be considered infinite) buffer size. We have also highlighted the complexity of analysing a network right after the stability limit as analytic models that consider the decoupling approximation can provide two different solutions and simulations have to be run for a long time to allow the system to move to its long-term behaviour. Moreover, we have also provided techniques to obtain the long-term performance, both analytical and experimentally. Computing the maximum service rate the network can support allows us to identify the range of packet arrival rates for which a potential for two solutions exist. We have also suggested how to set the starting parameters of iterative numerical solvers to obtain both solutions. Experimentally, we have shown that the transitory phase can be drastically ameliorated if stations begin with a number of packets preloaded in their buffers. Indeed, this is a useful technique to identify MAC protocols which may be subject to a transitory phase via simulation.

The potential occurrence of the extremely long transitory phase in generic random access protocols, as well as the characterisation of its duration, are important aspects to be analysed. They can provide further insights into the behaviour of networks based on random access protocols and serve as practical information to perform experimental evaluations. However, this study has been considered out of the scope of this work due to the complexity of analysing a system of coupled queues with state-dependent service rates, the reader is referred to [30] instead.

We have also evaluated the effect of the deferral counter on the magnitude of misprediction errors and shown that, under the same conditions, protocols such as the DCF have the potential to be affected by higher misprediction errors than the Homeplug MAC. On the other side, improvements of Homeplug MAC that increase the saturation throughput as the number of nodes increases show smaller misprediction errors.

Acknowledgment

This work has been partially supported by the Science Foundation Ireland grant 08/SRC/I1403 and 07/SK/I1216a.

References

[1] H. C. Ferreira, L. Lampe, J. Newbury, T. G. Swart, Power line communications, John Wiley & Sons Ltd, 2010.

[2] Homeplug Powerline Alliance, Homeplug 1.0 Specification.

[3] IEEE Std 1901, Standard for Broadband over Power Line Networks: Medium Access Control and Physical Layer Specifications, ANSI/IEEE Std 1901.

[4] IEEE Std 802.11, Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications, ANSI/IEEE Std 802.11.

[5] P. B. Velloso, M. Elias, M. Campista, D. D. O. Cunha, L. H. M. K. Costa, O. C. M. B. Duarte, Analyzing the performance of wireless local area networks with an improved collision avoidance mechanism, Technical Report.

[6] L. Jian, S. Qiang, The study on the performance of backoff algorithms in multihop power line communication networks, in: Proceedings of the Third International Conference on Measuring Technology and Mechatronics Automation (ICMTMA), Vol. 3, 2011, pp. 974–977.

[7] C. Vlachou, J. Herzen, P. Thiran, Fairness of MAC protocols: IEEE 1901 vs. 802.11, in: Proceedings of the IEEE International Symposium on Power Line Communications and its Applications (ISPLC), 2013.

[8] C. Cano, D. Malone, Evaluation of the backoff procedure of Homeplug MAC vs. DCF, in: Proceedings of the IEEE International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC 2013), 2013.

[9] M. Y. Chung, M.-H. Jung, T.-J. Lee, Y. Lee, Performance analysis of HomePlug 1.0 MAC with CSMA/CA, IEEE Journal on Selected Areas in Communications 24 (7) (2006) 1411–1420.

[10] S.-G. Yoon, J. Yun, S. Bank, Adaptive contention window mechanism for enhancing throughput in HomePlug AV networks, in: Proceedings of the 5th IEEE Consumer Communications and Networking Conference (CCNC), 2008, pp. 190–194.

[11] M. Koutny, P. Mlynek, P. Mrakava, Homeplug simulation model for analysis of the rounding functions, in: Proceedings of the 34th International Conference on Telecommunications and Signal Processing (TSP), 2011, pp. 218–221.

[12] P. J. Piñero-Escuer, J. Malgosa-Sanahuja, P. Manzanares-Lopez, J. P. Muñoz-Gea, Homeplug-AV CSMA/CA Evaluation in a Real In-Building Scenario, IEEE Communications Letters 15 (6) (2011) 683–685.

[13] M. Koutny, P. Mlynek, J. Misurec, Analysis of CSMA/CA used in power line communication, in: Proceedings of the 36th International Conference on Telecommunications and Signal Processing (TSP), 2013, pp. 385–388.

[14] P. J. Piñero, J. A. Cortés, F. J. Canete, L. Diez, P. Manzanares, J. Malgosa, A Realistic HomePlug-AV Simulator for In-home Network Services Planning, in: Proceedings of the Global Telecommunications Conference (GLOBECOM 2011), 2011, pp. 1–5.

[15] S.-G. Yoon, D. Kang, S. Bahk, Multichannel CSMA/CA Protocol for OFDMA-Based Broadband Power-Line Communications, IEEE Transactions on Power Delivery 28 (2013) 2491–2499.

[16] A. Kumar, E. Altman, D. Miorandi, M. Goyal, New insights from a fixed point analysis of single cell IEEE 802.11 WLANs, in: Proceedings of the IEEE International Conference on Computer Communications (INFOCOM), Vol. 3, 2005, pp. 1550–1561.

[17] G. Bianchi, I. Tinnirello, Remarks on IEEE 802.11 DCF performance analysis, IEEE Communication Letters 9 (8) (2005) 765–767.

[18] M. E. M. Campista, L. H. M. K. Costa, O. C. M. B. Duarte, Improving the data transmission throughput over the home electrical wiring, in: Proceedings of the IEEE Conference on Local Computer Networks, 2005. 30th Anniversary., 2005, pp. 318–327.

[19] C. Vlachou, A. Banchs, J. Herzen, P. Thiran, Performance Analysis of MAC for Power-Line Communications, in: Proceedings of the ACM SIGMETRICS Conference, 2014.

[20] C. Vlachou, A. Banchs, J. Herzen, P. Thiran, On the MAC for Power-Line Communications: Modeling Assumptions and Performance Tradeoffs, in: Technical Report, 2014.

[21] E. Kriminger, H. Latchman, Markov chain model of homeplug CSMA MAC for determining optimal fixed contention window size, in: Proceedings of the 2011 IEEE International Symposium on Power Line Communications and its Applications (ISPLC), 2011, pp. 399–404.

[22] K. Medepalli, F. A. Tobagi, System centric and user centric queueing models for IEEE 802.11 based wireless LANs, in: Proceedings of the 2nd International Conference on Broadband Networks, 2005. BroadNets 2005, 2005, pp. 612–621.

[23] B. Bellalta, M. Oliver, M. Meo, M. Guerrero, A simple model of the IEEE 802.11 MAC protocol with heterogeneous traffic flows, in: Proceedings of the IEEE Eurocon, 2005.

[24] K. R. Duffy, Mean field Markov models of wireless local area networks, Markov Processes and Related Fields 16 (2) (2010) 295–328.

[25] K. D. Huang, K. R. Duffy, D. Malone, On the validity of IEEE 802.11 MAC modeling hypotheses, IEEE/ACM Transactions on Networking 18 (6) (2010) 1935–1948.

[26] S. Borst, M. Jonckheere, L. Leskelä, Stability of parallel queueing systems with coupled service rates, Discrete Event Dynamic Systems 18 (4) (2008) 447–472.

[27] K. H. Suleiman, T. Javidi, M. Liu, S. Kittipiyakul, The impact of MAC buffer size on the throughput performance of IEEE 802.11, Technical Report.

[28] G. Chen, J. Branch, M. Pflug, L. Zhu, B. Szymanski, SENSE: A Sensor Network Simulator, Advances in Pervasive Computing and Networking (2004) 249–267.

[29] D. Malone, K. R. Duffy, D. Leith, Modeling the 802.11 distributed coordination function in nonsaturated heterogeneous conditions, IEEE/ACM Transactions on Networking 15 (1) (2007) 159–172.

[30] C. Cano, D. Malone, Modeling, Analysis and Impact of a Long Transitory Phase in Random A
abs/1405.5753.
URL http://arxiv.org/abs/1405.5753

Cristina Cano obtained the Telecommunications Engineering Degree at the Universitat Politecnica de Catalunya (UPC) in February 2006. Then, she received her M.Sc. (2007) and Ph.D. (2011) on Information, Communication and Audiovisual Media Technologies from the Universitat Pompeu Fabra (UPF). Since July 2012, she has been working as a research fellow at the Hamilton Institute (NUIM) in wireless networks, sensor networks, power line communications and MAC layer design.

David Malone received B.A. (mod), M.Sc. and Ph.D. degrees in mathematics from Trinity College Dublin. During his time as a postgraduate, he became a member of the FreeBSD development team. He is currently a SFI Stokes Lecturer at the Hamilton Institute, NUI Maynooth. His interests include mathematics of networks, network measurement, IPv6 and systems administration. He is a coauthor of *O'Reilly's IPv6 Network Administration*.